

IFW



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/811,410	03/26/2004	Toshihiko Kataoka	JP920030050US1	3276

53493 7590 09/27/2005

LENOVO (SINGAPORE) PTE. LTD.
BUILDING 675, MAIL C-137
4401 SILICON DRIVE
DURHAM, NC 27709

EXAMINER

ZAMAN, FAISAL M

ART UNIT PAPER NUMBER

2112

DATE MAILED: 09/27/2005

RECEIVED
OICE/IAP

NOV 01 2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/811,410	KATAOKA, TOSHIHIKO	
	Examiner	Art Unit	
	Faisal Zaman	2112	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 20 July 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-17 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-17 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☒ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 26 March 2004 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☒ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☒ All b) ☐ Some * c) ☐ None of:
1. ☒ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| Paper No(s)/Mail Date <u>July 20, 2004</u> . | 6) <input type="checkbox"/> Other: _____ |

PD

DETAILED ACTION

Information Disclosure Statement

1. The references listed on the Information Disclosure Statement submitted on 22 July 2004 have been considered by the examiner (see attached PTO-1449).

Specification

2. The title of the invention is not descriptive. A new title is required that is clearly indicative of the invention to which the claims are directed.

The following title is suggested: --INTERRUPT CONTROL DEVICE WHICH SENDS DATA TO PROCESSOR AT OPTIMIZED TIME--.

In addition, please make the title the same throughout all documents in the application ("Method for Data Protection For Removable Recording Medium" was used as the title in several documents in the application).

3. The disclosure is objected to because of the following informalities:
 - a. In Line 6 of the Background, the Published Unexamined Patent Application No. should be JP10-275136, rather than 10-275136.
 - b. In Line 4 of the third paragraph of the Background, "polling" is misspelled.Appropriate correction is required.

Claim Rejections - 35 USC § 112

4. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Art Unit: 2112

5. **Claims 1-17** are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

As per Claims 1, 15, 16, and 17, if the interrupt is issued prior to data acquisition, it is not clear how the interrupt indicates the data has become available. The Examiner would interpret this limitation, for examination purposes, to mean the interrupt is sent to the central processing unit before said object acquiring unit receives *all* said data or said resource, indicating that *some* said data or said resource has become available.

As per Claim 7, it is not clear as to what a “predetermined small value” is or how the “predetermined small value” is determined. The Examiner would interpret this limitation, for examination purposes, to mean the setup period change unit changes said setup period to make said average the *smallest possible* value.

As per Claim 8, it is not clear as to what type of “distribution of said time differences” the setup period change unit changes the setup period according to. The Examiner would interpret this limitation, for examination purposes, to mean the setup period change unit changes said setup period according *to an average* of a distribution of said time differences measured by said time difference measuring unit.

As per Claim 9, it is not clear as to what a “predetermined percentage of said time differences” is or how it is determined. In addition, it is not clear as to what the “predetermined value” is or how it is determined. The Examiner would

Art Unit: 2112

interpret this limitation, for examination purposes, to mean to have said setup period change unit change said setup period to make the *average* of time differences less than or equal to the *smallest possible* value. The "predetermined percentage of said time differences" would be the said time differences in the average which are below the smallest possible value described in Claim 7.

All claims not specifically addressed are rejected due to a dependency.

Correction/clarification is therefore required.

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1, 15, 16, and 17 are rejected under 35 USC 103(a) as being obvious over Hashimoto et al. ("Hashimoto") (U.S. 6,397,282) in view of Williams et al. ("Williams") (U.S. 6,167,480).

Hashimoto discloses the invention substantially as claimed.

Regarding Claim 1, Hashimoto discloses:

An interrupt control device (Fig. 3, item 20, Column 2, lines 23-29) for issuing interrupts to a central processing unit, comprising:

An object acquiring unit (Fig. 3, item 22, Column 11, line 1) for acquiring data or resource(s) for use by said central processing unit;

An interrupt issuing unit for issuing an interrupt to said central processing unit before said object acquiring unit acquires said data or said resource, said interrupt indicating that said data or said resource has become available (Fig. 3, item 26, Column 11, lines 15-19);

Hashimoto does not disclose expressly:

A use delay unit for delaying the use of said data or said resource by said central processing unit until said object acquiring unit acquires said data or said resource if said central processing unit which has received said interrupt requests the use of said data or said resource before said object acquiring unit acquires said data or said resource.

In the same field of endeavor (e.g. providing an interrupt signal to a host system processor), Williams discloses a use delay unit (Column 3, lines 8-21) for delaying the use of said data or said resource by said central processing unit until said object acquiring unit acquires said data or said resource if said central processing unit which has received said interrupt requests the use of said data or said resource before said object acquiring unit acquires said data or said resource.

Accordingly, it would have been obvious to one of ordinary skill in the art at the time of the invention was made to have incorporated Williams' teachings of providing an interrupt signal to a host system processor with the teachings of Hashimoto, for the purpose of providing only a single interrupt for an incoming information packet's reception (see Williams, Column 4, lines 8-10). Also, it would have been desirable as stated by Williams for a network peripheral to be

Art Unit: 2112

able to provide such an interrupt for information packets that conform to a user specified network operating system protocol (see Williams, Column 4, lines 8-26). Hashimoto provides motivation to combine by making a point of his invention to be able to postpone interrupt requests that do not require immediate attention (see Hashimoto, Column 2, lines 10-22).

Regarding Claim 15 all the same elements of Claim 1 are listed, but where the central processing unit and interrupt control device are located in an information processing device. Since it would be obvious to one of ordinary skill in the art that a central processing unit and interrupt control device would be located in some sort of an information processing device, the supporting rationale of the rejection to Claim 1 applies equally as well to Claim 15.

Regarding Claim 16 all the same elements of Claim 1 are listed, but in method form rather than system form. Therefore, the supporting rationale of the rejection to Claim 1 applies equally as well to Claim 16.

Regarding Claim 17 all the same elements of Claim 1 are listed, but in program product form rather than system form. Therefore, the supporting rationale of the rejection to Claim 1 applies equally as well to Claim 17.

Claim Rejections - 35 USC § 103

7. **Claims 2-5, 10, and 12-14** are rejected under 35 U.S.C. 103(a) as being unpatentable over Hashimoto-Williams as applied to claim 1 above, and further in view of Reid et al. ("Reid") (U.S. 6,115,776).

Hashimoto-Williams discloses the invention substantially as claimed.

Art Unit: 2112

Hashimoto-Williams discloses the interrupt control device according to Claim 1 as described above.

Regarding Claim 2, Hashimoto-Williams does not disclose expressly wherein said interrupt issuing unit issues said interrupt after a predetermined setup period elapses from when a data generation device generating said data starts to generate said data.

In the same field of endeavor (e.g. improvements to the transmission of information between digital devices over a communications medium), Reid discloses wherein said interrupt issuing unit (Reid, Column 9, lines 6-15) issues said interrupt after a predetermined setup period elapses from when a data generation device generating said data starts to generate said data.

Accordingly, it would have been obvious to one of ordinary skill in the art at the time the invention was made to have incorporated Reid's teachings of improvements to the transmission of information between digital devices over a communications medium with the teachings of Hashimoto-Williams, for the purpose of reducing "the number of interrupts generated by the network to the processor and to thereby reduce the processing burden to the operating system of servicing those interrupts" (Reid, Column 3, lines 29-31). Hashimoto-Williams provides motivation to combine by stating that interruption requests that do not require immediate attention are postponed from being sent to the central processing unit until a predetermined delay has elapsed (Hashimoto, Column 2, lines 10-22).

Art Unit: 2112

Regarding Claim 3, Hashimoto-Williams and Reid disclose the following limitations, which are not disclosed expressly in Hashimoto and Williams:

A time difference measuring unit (Reid, Column 9, lines 6-15) for measuring a time difference between when said object acquiring unit acquires said data and when said central processing unit which has received said interrupt requests the use of said data; and

A setup period change unit (Column 9, lines 27-29) for changing said predetermined setup period according to said time difference.

The motivation that was utilized in the combination of Claim 2, super, applies equally as well to Claim 3.

Regarding Claim 4, Hashimoto and Williams do not disclose expressly:

An acquisition time measuring unit for measuring an acquisition time from when said data generation device starts to generate said data until said object acquiring device acquires said data;

Wherein said setup period change unit changes said setup period according to said acquisition time and said time difference.

In the same field of endeavor, Williams discloses:

An acquisition time measuring unit (Figure 11, item 270, Column 14, lines 14-18) for measuring an acquisition time from when said data generation device starts to generate said data until said object acquiring device acquires said data;

Wherein said setup period change unit changes said setup period according to said acquisition time and said time difference.

Art Unit: 2112

Please note the definition of "latency" as used in the reference in this rejection can be found at <Hyperdictionary.com>. The portion of the claim regarding the setup period change unit is rejected because it would be obvious to one of ordinary skill in the art that said acquisition time and the said time difference are directly related, therefore the supporting rationale of the rejection to Claim 3 applies equally as well here.

The motivation that was utilized in the combination of Claim 2, super, applies equally as well to Claim 4.

Regarding Claim 5, Reid discloses the following limitations, which are not disclosed expressly in Hashimoto and Williams:

Said data generation device (Column 3, lines 40-44) generates a plurality of data segments;

Said object acquiring unit (Column 5, lines 3-6) sequentially acquires said plurality of data segments for use by said central processing unit;

Said interrupt issuing unit (Column 9, lines 7-15) issues an interrupt to said central processing unit before said object acquiring unit acquires each of said plurality of data segments, each said interrupt indicating that the respective one of said plurality of data segments has become available;

Said time difference measuring unit (Column 9, lines 6-15) measures, for each of said plurality of data segments, the time difference between when said object acquiring unit acquires said data segment and when said central processing unit which has received said interrupt requests the use of said data segment; and

Said setup period change unit (Column 9, lines 27-29) changes said setup period according to the time differences measured by said time difference measuring unit.

The motivation that was utilized in the combination of Claim 2, super, applies equally as well to Claim 5.

Regarding Claim 10, Hashimoto and Williams do not disclose a setup period change unit for, (i) changing said setup period to a smaller value if said central processing unit which has received said interrupt requests the use of said data or said resource before said object acquiring unit acquires said data or said resource, and (ii) changing said setup period to a greater value if said central processing unit which has received said interrupt requests the use of said data or said resource after said object acquiring unit acquires said data or said resource. However, in the same field of endeavor, Reid discloses a setup period change unit which has a delay that may be user-programmable, or it may be system-programmable based on varying system parameters (Reid, Column 6, lines 36-37). Reid further discloses a setup period change unit that is able to generate an interrupt after a said number of packets accumulates or after a predetermined period of time elapses (Reid, Column 9, lines 13-15).

Accordingly, it would have been obvious to one of ordinary skill in the art at the time of the invention was made to have incorporated Reid's teachings of delaying interrupts until a predetermined number of data packets accumulates or a predetermined period of time elapses to the teachings of Hashimoto-Williams for the purpose of increasing or decreasing the number of data packets that are

Art Unit: 2112

stored in storage (Hashimoto, Column 11, line 1) before an interrupt is sent to the central processing unit.

The motivation that was utilized in the combination of Claim 2, super, applies equally as well to Claim 10.

Regarding Claims 12, 13, and 14, all the same elements of Claims 2, 3, and 4, respectively are listed, but where “data generation device” is replaced with “resource reservation device”. Since “data” and “resource” were used interchangeably in other claims (ie. Claims 1, 10, and 11, and numerous times in the specification), the supporting rationale of the rejection to Claims 2, 3 and 4 apply equally as well to Claims 12, 13, and 14, respectively.

Claim Rejections - 35 USC § 103

8. **Claims 6-9** are rejected under 35 USC 103(a) as being unpatentable over Hashimoto-Williams and Reid as applied to Claims 2-5 above, and in further view of Williams (“Williams '305”) (U.S. 6,061,305).

Hashimoto-Williams and Reid disclose the invention substantially as claimed.

Regarding Claim 6, Hashimoto-Williams do not disclose expressly wherein said setup period change unit changes said setup period according to the average of the time differences measured by said time difference measuring unit.

In the same field of endeavor (e.g. providing an interrupt signal to a host processor unit), Hashimoto-Williams and Reid disclose wherein said setup period change unit (Williams, Column 14, lines 19-26) changes said setup period according to the average of the time differences (Williams '305, Column 4 line 66 – Column 5, line 27) measured by said time difference measuring unit.

Art Unit: 2112

Accordingly, it would have been obvious to one of ordinary skill in the art at the time the invention was made to have incorporated Williams '305's teachings of providing an interrupt signal to a host processor unit along with measuring duration times of each of a series of events and particularly to determining the average event duration time to the teachings of Hashimoto-Williams and Reid, for the purpose of finding a value that will be most useful in optimizing device performance (Williams '305, Column 1, lines 39-40).

Hashimoto-Williams and Reid provide motivation to combine by stating that an interrupt is asserted as an optimum interrupt time (Williams, abstract).

Regarding Claim 7, Williams '305 discloses the following limitation, which is not disclosed expressly in Hashimoto-Williams: wherein said setup period change unit (Williams '305, Column 5, lines 23-27) changes said setup period to make said average a predetermined small value.

The motivation that was utilized in the combination of Claim 6, super, applies equally as well to Claim 7.

Regarding Claims 8 and 9, since they are directly related to Claim 7 (according to the Examiner's interpretation), the supporting rationale of the rejection to Claim 7 applies equally as well to Claims 8 and 9.

Claim Rejections - 35 USC § 103

9. **Claim 11** is rejected under 35 USC 103(a) as being obvious over Hashimoto-Williams in view of Reid and in further view of Brice, Jr. et al. ("Brice") (U.S. 6,754,738).

Hashimoto-Williams discloses the invention substantially as claimed.

Hashimoto-Williams discloses a delay processing unit for (ii) causing said central processing unit to return from interrupt handling caused by said interrupt, to delay the use of said data or said resource by said central processing unit until said object acquiring unit acquires said data or said resource (Hashimoto, Figure 5, Column 8, lines 1-24). However, Hashimoto-Williams do not disclose expressly a delay time calculation unit for calculating a delay time required from the time said object acquiring unit receives a request for use of said data or resource from said central processing unit which has received said interrupt until said object acquiring unit acquires said data or resource as well as a delay processing unit for (i) causing said central processing unit to use polling to request said data or said resource if said delay time is less than a predetermined threshold.

In the same field of endeavor (e.g. improvements to the transmission of information between digital devices over a communications medium), Reid discloses a delay time calculation unit (Column 10, lines 8-18) for calculating a delay time required from the time said object acquiring unit receives a request for use of said data or resource from said central processing unit which has received said interrupt until said object acquiring unit acquires said data or resource.

Accordingly, it would have been obvious to one of ordinary skill in the computer architecture art at the time the invention was made to have incorporated Reid's teachings of a computer operating system in which interrupts are generated to a processor by events which then require processor time to service to the teachings of Hashimoto and Williams, for the purpose of delaying

Art Unit: 2112

data to be sent to the processor until a predetermined amount of data segments or a predetermined time has elapsed (Reid, abstract, Brice, abstract).

Hashimoto-Williams provides motivation to combine by stating that interruption requests that do not require immediate attention are postponed from being sent to the central processing unit a predetermined delay has elapsed (Hashimoto, Column 2, lines 10-22).

In the same field of endeavor, Brice discloses a delay processing unit (Brice, Column 11, lines 28-32) for (i) causing said central processing unit to use polling to request said data or said resource if said delay time is less than a predetermined threshold.

The motivation that was utilized in the combination of the previous part of the claim, super, applies equally as well to this part of the claim.

Conclusion

10. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Jinzaki (U.S. Publication No. 2004/0236875) discloses a computer for determining interruption delay dynamically. Williams et al. (U.S. 5,881,296) discloses a method for improved interrupt processing in a computer system. Huffman et al. (U.S. 6,640,274) discloses a method and apparatus for reducing the disk drive data transfer interrupt service latency penalty. Paul et al. (U.S. 6,721,878) discloses low-latency interrupt handling during memory access delay periods in microprocessors. Stevens (U.S. 6,338,111) discloses a method and apparatus for reducing I/O interrupts. Kailash et al. (U.S. 6,185,639) discloses a system and method to reduce a computer system's interrupt

Art Unit: 2112

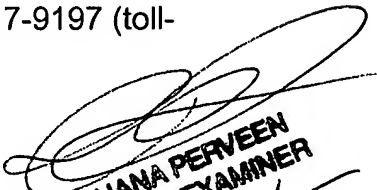
processing overhead. Binford et al. (U.S. 5,671,365) discloses an I/O system for reducing main processor overhead in initiating I/O requests and servicing I/O completion events. Bashford (U.S. 6,629,179) discloses a message signaled interrupt generating device and method. Constantinos Dovrolis, Brad Thayer, and Parameswaran Ramanathan (ACM SIGOPS Operating Systems Review, Volume 35, Issue 4) disclose a method for hybrid interrupt-polling for the network interface.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Faisal Zaman whose telephone number is 571-272-6459. The examiner can normally be reached on Monday thru Friday, 9 am – 5:30 pm.


If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Rehana Perveen can be reached on 571-272-3676. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

FMZ


REHANA PERVEEN
PRIMARY EXAMINER
9/26/05

FORM PTO - 1449 (Modified) LIST OF PATENTS AND PUBLICATIONS FOR APPLICANT'S INFORMATION DISCLOSURE STATEMENT (Use several sheets if necessary) Sheet 1 of 1	Application Number	10/811,410
	Filing Date	March 26, 2004
	First Named Inventor	Toshihiko Kataoka
	Group Art Unit	2189 2112
	Confirmation No.	3276
	Attorney Docket Number	JP920030050US1



UNITED STATES PATENT DOCUMENTS						
Examiner Initials	Cite No.	Patent Document Number	Kind Code	Inventor	Date of Publication mm/dd/yyyy	Pages, Columns, Lines Where Relevant Passages Appear

FOREIGN PATENT DOCUMENTS						
Examiner Initials	Cite No.	Patent Document Number	Kind Code	Country	Date of Publication mm/dd/yyyy	Pages, Columns, Lines Where Relevant Passages Appear
FMZ	F1	JP10-275136		JAPAN	10/13/1998	ABSTRACT ONLY

OTHER ART (Including Author (CAPITAL LETTERS), Title, Date, Pertinent Pages, etc.)		
Examiner Initials	Cite No.	

Examiner Signature 	Date Considered 9/14/05
-------------------------------	-----------------------------------

EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

Notice of References Cited	Application/Control No. 10/811,410		Applicant(s)/Patent Under Reexamination KATAOKA, TOSHIHIKO	
	Examiner Faisal Zaman		Art Unit 2112	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
	A	US-6,397,282	05-2002	Hashimoto et al.	710/260
	B	US-6,167,480	12-2000	Williams et al.	710/260
	C	US-6,115,776	09-2000	Reid et al.	710/260
	D	US-6,061,305	05-2000	Williams, Robert Alan	368/113
	E	US-6,754,738	06-2004	Brice et al.	710/48
	F	US-2004/0236875	11-2004	Jinzaki, Akira	710/015
	G	US-5,881,296	03-1999	Williams et al.	710/263
	H	US-6,640,274	10-2003	Huffman et al.	710/260
	I	US-6,721,878	04-2004	Paul et al.	712/244
	J	US-6,338,111	01-2002	Stevens, Jerry Wayne	710/260
	K	US-6,185,639	02-2001	Kailash et al.	710/48
	L	US-5,671,365	09-1997	Binford et al.	710/100
	M	US-6,629,179	09-2003	Bashford, Patrick R.	710/260

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Constantinos Dovrolis, Brad Thayer, Parameswaran Ramanathan, "HIP: Hybrid Interrupt-Polling for the Network Interface", 2001, ACM SIGOPS Operating Systems Review, ACM Press, Volume 35, Issue 4, Pages 50-60.
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

HIP: Hybrid Interrupt-Polling for the Network Interface

Constantinos Dovrolis¹ Brad Thayer² Parameswaran Ramanathan³

¹Department of Computer and Information Sciences, University of Delaware

²Department of Computer Sciences, University of Wisconsin

²Department of Electrical and Computer Engineering, University of Wisconsin
dovrolis@cis.udel.edu, brad@cs.wisc.edu, parmash@ece.wisc.edu

Abstract

The standard way to notify the processor of a network event, such as the arrival or transmission of a packet, is through interrupts. Interrupts are more effective than polling, in terms of the per packet send/receive latency. Interrupts, however, incur a high overhead both during and after the interrupt handling, because modern superscalar processors use long pipelines, out-of-order and speculative execution, and multi-level memory systems, all of which tend to increase the interrupt overhead in terms of clock cycles. In this paper, we attempt to reduce the network interface overhead by introducing a hybrid scheme (HIP) that uses interrupts under low network load conditions and polling otherwise. Even though such hybrid schemes have been proposed in the past, the polling period in HIP is adjusted dynamically based on the rate of the arriving packet stream. In this way, the increase in the per packet latency, which occurs with polling, is quite low. This is quantified with trace-driven simulations, which also show that the per packet overhead with HIP is significantly reduced compared to the conventional interrupt-based mechanism. HIP would be beneficial for high bandwidth network interfacing in servers with a heavy WWW or streaming media workload.

1 Introduction

Personal computers and workstations will soon be used as videophones, televisions, and multiplayer game systems. These applications are network-intensive, and due to their multimedia nature they require high-throughput network interfaces. Many workstations today are connected to Fast Ethernet interfaces (100Mbps), while Gigabit Ethernet interfaces (1000Mbps) are already deployed in high-end Web

servers. This dramatic bandwidth increase calls for optimizations in all key components of the network interface, including the network interface card (NIC), the protocol stacks, the operating system, the input-output unit, the memory system, and the processor. In general, the network interface has been traditionally viewed as just an I/O peripheral that causes unpredictable and infrequent events, and so not much optimization has been put into it [1, 2, 3, 4, 5].

We first have to differentiate between the concepts of latency and throughput, as they apply to the receive operation from the network interface. *Latency* is the time duration between a packet arrival at the NIC and its delivery to the application. The *throughput*, on the other hand, is the rate with which the application can read packets from the NIC. The throughput is the inverse of the receive overhead, where the latter accounts for all processing delays of the receive operation. The latency can be larger than the overhead if the received packets are queued at the NIC before they are received and processed by the application. There are systems and applications where reducing latency is the major issue. For example, in a loosely-coupled multiprocessor system the network latency is part of the computational delay and, therefore, it has to be minimized [6]. In contrast, in multimedia applications the audio and video streams that arrive from the network are queued for a playback delay of several tens or hundreds of milliseconds before they are delivered to the user [7]. Consequently, an increase of a few milliseconds in the receive latency would normally go unnoticed from that type of applications.

Current personal computers and workstations are notified of an arrival of a packet at the NIC through interrupts. This is because interrupts guarantee a minimum receive latency since the packets are not queued at the NIC for a duration more than the interrupt handling period. An interrupt, however, is an asynchronous event with high hardware and soft-

ware overhead. The hardware overhead is mainly due to the flushing of out-of-order and speculative execution state in the processor and due to the reduction in the locality of references in the instruction and data caches caused by the context switches that are necessary for interrupt processing [8]. The software overhead is due to the following reasons. When an interrupt occurs, the architecturally visible state must be saved, an appropriated interrupt handler must be dispatched, and upon completion of the handler, the system state must be restored. In addition, to handle nested interrupts, appropriate process states and priorities must be updated in each invocation of the interrupt handler, requiring time-consuming bookkeeping [9].

An alternative to interrupts is polling. In polling, the processor periodically initiates a read operation of a control NIC register. If one or more packets have arrived, they are moved to the main memory for further processing. Since several packets may be read in the same poll and since the code to perform a poll is usually much shorter than the interrupt processing code, the receive overhead is reduced. However, on the negative side, packets are not guaranteed to be present at each poll; the polls in which no packet is found in the NIC (unsuccessful polls) increase the overall overhead of the network interface. Additionally, the latency of the receive operation increases because packets are queued in the NIC until the polling event. Because of these two drawbacks, polling is not commonly used in general purpose systems. Polling is used however in systems that have a heavy network load, such as routers and bridges, firewalls, or file servers [5]. In such systems the probability of unsuccessful polls is small, and the receive latency remains quite low by using a high polling period, and/or specialized hardware (see Section 5).

In this paper, we propose an input network interface mechanism which combines the advantages of interrupts and polling, i.e., it has a receive overhead that is comparable to that of polling, and a receive latency that is not prohibitively larger than that of interrupts. This scheme is based on the following two observations about next generation multimedia applications.

1. Multimedia network traffic is stream-based with frequent packet arrivals that have some statistical predictability. By monitoring the packet interarrival times, the next packet arrival time

can be roughly estimated. If this prediction is effective, we can in principle set the polling period to the estimated packet interarrival period, reducing both the unsuccessful polls as well as the receive latency. Notice that in all current implementations of polling-based network interfaces, the polling period is fixed and independent of the arriving packet stream interarrivals.

2. Multimedia-based applications can tolerate an increase in the receive latency, if this increase is much smaller than the playback delay of the application. The playback delay is the time duration between the receipt of a packet and the display of its data to the user. The playback delay is typically used for absorbing the random delays in the network. For interactive real-time applications, the playback delay is in the order of 100–200 milliseconds. Consequently, an increase in the receive latency of a few milliseconds should not be a problem for such applications. For other applications, such as Remote Procedure Calls (RPC) or Network File System (NFS) operations, the latency should not be larger than, for example, the average latency of a disk operation. Since a disk access can be several milliseconds, a network interface latency of the same magnitude should not cause noticeable performance degradations.

The basic idea of the proposed network interface scheme, called *Hybrid Interrupt Polling* (HIP), is to adaptively switch between the use of interrupts and polling based on the observed rate of packet arrivals. Specifically, if the packet arrivals are frequent and predictable, the receive mechanism operates in a polling mode and the interrupts are disabled. In this mode, the polling period is set based on the predicted packet interarrival times. However, to bound the receive latency, the polling period is not allowed to exceed a pre-determined limit (say, ten milliseconds). On the other hand, if the packet arrivals are infrequent, less predictable, or if the number of consecutive unsuccessful polls exceeds a threshold, the receive mechanism operates in the interrupt mode. In this mode, the polling operation is stopped and the interrupts are enabled.

HIP performs better than interrupts in terms of receive overhead and better than polling in terms of receive latency. Additionally, the receive latency is always upper bounded, as a safety precaution of the

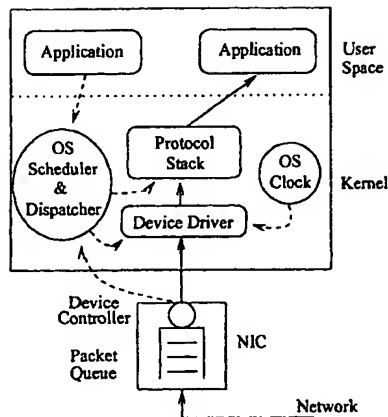


Figure 1: The basic components of the receive part of the network interface

mechanism. We demonstrate these two results by simulating HIP and other network interface mechanisms using traces from real local-area network traffic.

The rest of this paper is as follows. Section 2 presents a simple model for the network interface and it illustrates the main performance measures. Section 3 explains how polling is performed in HIP. Section 4 presents HIP and the involved algorithms in detail. Section 5 discusses some related works and their differences and similarities with HIP. Section 6 evaluates quantitatively HIP and other schemes through simulation. Finally, Section 7 concludes and identifies remaining open issues.

2 A Model of the Network Interface

The network interface is a complex system that consists of several hardware and software interacting units in both the computer and the NIC. We show the major components of the receive part of this system in Figure 1. As [4] points out, the receive and the transmit parts of the network interface are not completely symmetrical. In this paper we focus on the receive-part, where the interrupt overhead is more important.

We consider a typical multitasking workstation, where all the network interface functionality is performed by operating system processes running in the

kernel address space, while the application processes run in the user address space. When a packet arrives at the NIC it gets temporarily stored in a local queue. In an interrupt-based system, the NIC device controller can either transfer the packet in the kernel memory buffers using a DMA engine and then interrupt the processor, or it can immediately interrupt the processor that will then copy the packet to the kernel memory. Currently, most network interfaces are DMA-capable. The activation of the hardware interrupt line preempts the running process, and the OS interrupt dispatcher is invoked to identify the nature of the interrupt and the corresponding device driver. After an additional context switch, the device driver completes the receive transaction with the device controller, and the packet processing continues with the network protocol processes (e.g., IP, UDP). Finally, the packet is copied again, this time from the kernel space to the user space, and the recipient application is notified.

Some of the above operations have a fixed overhead per packet, while some others have an overhead that is roughly proportional to the packet size in bytes. Following the model of [4], the final throughput R that the application gets is

$$R = \frac{1}{V_P + L \times V_B}$$

where L is the packet length in bytes, V_P is the per-packet overhead, and V_B is the per-byte overhead. The fixed overhead accounts for the hardware cost of the DMA transfer setup and of the I/O bus arbitration, as well as for the execution of the interrupt service routine, the device driver, the protocol stack, and for costs associated with memory management and context switches. The per-packet overhead is mainly due to data copying from the NIC to the kernel space (if there is no DMA support), and from the kernel space to the user process, as well as with error-checking processing, such as checksum calculations. In this paper we attempt to optimize the component of the fixed overhead per packet that is related with the interrupt overheads. As measured in [10], a hardware interrupt (such as the interrupt generated by the NIC) with a null interrupt handler introduces an overhead of about $4 \mu s$ in a 500Mhz Pentium III system running FreeBSD 2.2.6. This is an important overhead for a workstation that receive streams of several hundreds of Mbps from the network interface, and it is important to examine possible optimizations that can reduce it.

3 The Polling Process

Before describing HIP in detail, we address the important issue of how polling is performed. The constraints that HIP imposes are, first, that the polling period has to be dynamically adjustable, so that it can track the packet interarrivals, and second, that the polling period has to have as fine a time granularity as possible, so that the packet latency to be minimized.

For a multitasking operating system, one possible polling approach is to have a periodically scheduled kernel polling process. This approach, however, requires a context switch for each poll. Although the overhead of this context switch is smaller than the cost of an interrupt (no hardware overhead and no interrupt dispatching), it is still comparable to the interrupt handling overhead. Ideally, the polling operation should not introduce a context switch overhead.

The adopted solution in HIP is based on the operating system soft clock [9]. This clock causes a periodic interrupt that is used for time-slicing and other bookkeeping activities. Its period is the finest time-slice and system clock granularity that the operating system allows. In HIP, the polling period is always set to a multiple of the clock period. Specifically, the next polling event is scheduled using a counter P in the clock handler. In every clock tick, the value of P gets decremented. When it reaches zero, the network device driver is called to poll the NIC. Then, the value of P is set to the number of clock events until the next poll, and the process repeats. This operation requires only a few additional instructions in the clock interrupt handler. The actual polling overhead is the cost of reading a status register on the NIC to check for a packet arrival, and if one is detected, to transfer the packet(s) from the NIC buffers to main memory (normally using DMA). Note that the overhead of the soft clock interrupt would be encountered in anyway, so the polling operation does not introduce an additional context switch.

Over the last decade or so, the OS clock period was commonly set to 10 milliseconds [9]. Although the polling period can be constrained to a multiple of this period, the time granularity of adjusting the polling period would be too coarse, and the maximum latency that polling could introduce would be excessive (several tens of milliseconds) for many ap-

plications. More recently however, some OS vendors have moved to a smaller clock interrupt period of 1 millisecond (e.g., Solaris 8). This reduction is well justified, given that the CPU performance has improved several orders of magnitude over the last ten years. We believe that most OS vendors will soon also switch to a 1-msec clock interrupt period, or even lower than that. As it will become clear in the next section, HIP becomes more practical and effective as the clock interrupt period decreases.

4 HIP

In this section we first present the important parameters in HIP, and then describe the algorithms to determine when to perform polling instead of interrupts, and how to compute the polling period.

- **Interrupt Overhead V_I :** The fixed overhead of receiving a packet from the network interface using interrupts.
- **Polling Overhead V_P :** The fixed overhead of receiving a packet from the network interface using polling. A rule of thumb is that the polling overhead is an order of magnitude less than the interrupt overhead [8, 6, 11].
- **Transfer Overhead $V(B)$:** The overhead of transferring B bytes from the network interface to main memory. This accounts for the variable part of the interrupt and polling overhead, i.e., if B bytes have arrived, the total interrupt overhead is $V_I + V(B)$, and the total polling overhead is $V_P + V(B)$. $V(B)$ is practically zero for a DMA-capable network interface. For simplicity, we assume that this overhead is proportional to B ($V(B) \propto B$).
- **Mode Switching Overhead V_S :** The overhead of switching from Interrupt-mode to Polling-mode, and vice versa. This overhead accounts for the network interface interrupt enabling and disabling operations. We assume that both operations have the same cost.
- **Minimum Polling Period T_P^{MIN} :** The minimum polling period that can be set. It is limited by the soft clock period.

- **Maximum Polling Period T_P^{MAX} :** The maximum polling period that can be set. It is limited by the maximum latency that can be tolerated by latency-sensitive applications such as NFS or RPC.
- **Polling Period T_P :** The scheduled polling period. The constraints are that $T_P^{MIN} \leq T_P \leq T_P^{MAX}$, and that T_P must be an integral multiple of T_P^{MIN} .
- **Average Packet Interarrival \bar{I} :** A dynamic estimation of the average duration between packet arrivals.
- **Last Packet Interarrival D :** The duration between the last packet arrival and the one before that.
- **Packet Interarrival Variance σ_I^2 :** A dynamic estimation of the packet interarrival variance. It is used as a rough statistical indication of the regularity of the packet interarrivals. The square root of σ_I^2 is the packet interarrival standard deviation σ_I .
- **Consecutive Unsuccessful Polls P_U :** The number of consecutive unsuccessful polls (polls where there was nothing received at the network interface). An unsuccessful poll increases the overhead without producing any useful work. A large value of P_U indicates that the polling period used was poorly set, or that the stream of received packets has ended.
- **Maximum Consecutive Unsuccessful Polls P_U^{MAX} :** The maximum allowed value of P_U . After this number of unsuccessful polls, HIP switches back to Interrupt-mode. As a first guess, P_U can be equal to the ratio $\frac{V_I}{V_P}$, since this many unsuccessful polls are equivalent, in terms of overhead, to an interrupt.

The average packet interarrival \bar{I} is estimated after each packet arrival using a simple exponential weighted average:

$$\bar{I} = \alpha \bar{I} + (1 - \alpha)D \quad (0 \leq \alpha \leq 1) \quad (1)$$

where α is the **average interarrival weight factor**. α controls the importance that is given to the last interarrival relative to the past history of interarrivals, as that is accumulated in \bar{I} .

The packet interarrival variance σ_I^2 is estimated as

$$\sigma_I^2 = \beta \sigma_I^2 + (1 - \beta)(\bar{I} - D)^2 \quad (0 \leq \beta \leq 1) \quad (2)$$

where β is the **interarrival variance weight factor**. Again, β controls the memory of the predictor.

HIP switches between the interrupt and the polling modes based on the observed statistics of packet arrivals. It switches to the polling mode if the following two conditions are satisfied:

1. Packets arrive regularly enough
2. Packets arrive frequently enough

The first condition attempts to avoid the risk of many unsuccessful polls when the packet stream is too bursty. The second condition avoids polling when the arrival rate is small compared to the minimum possible polling rate. Specifically, HIP switches to polling mode if

$$\frac{\sigma_I}{\bar{I}} < \gamma \quad \text{and} \quad \bar{I} < T_P^{MAX} \quad (3)$$

The ratio σ_I/\bar{I} is the coefficient of variation of the packet interarrivals and it is an indication of the predictability of the packet interarrivals. The parameter γ is the **prediction threshold**. Reducing γ requires less predictable packet interarrivals before switching to polling mode.

HIP switches to interrupt mode when the complement of the above conditions occur, or when the number of consecutive unsuccessful polls has reached its maximum allowed value. This latter condition ensures that when a stream has finished, or when the interarrivals prediction has failed, HIP will switch back to interrupts. Specifically, HIP switches to the interrupt mode if

$$\frac{\sigma_I}{\bar{I}} \geq \gamma \quad \text{or} \quad \bar{I} \geq T_P^{MAX} \quad \text{or} \quad P_U = P_U^{MAX} \quad (4)$$

When P_U reaches its maximum value P_U^{MAX} , it is re-initialized to zero, and \bar{I} and σ_I get their default values.

Ideally, when in polling mode, the polling period should be set to exactly match the next packet arrival instant. This, of course, is impossible, first because we cannot exactly predict future arrivals, and second, because the polling period has to be a multiple of T_P^{MIN} , and $T_P^{MIN} \leq T_P \leq T_P^{MAX}$. In HIP, the polling period is set based on the following algorithm.

$$\begin{aligned}
&T_P = \bar{I} + \phi\sigma_I; \\
&\text{If } (T_P < T_P^{MIN}) \\
&\quad T_P = T_P^{MIN}; \\
&\text{Else} \\
&\text{If } (T_P > T_P^{MAX}) \\
&\quad T_P = T_P^{MAX};
\end{aligned}$$

In this algorithm, ϕ is the **slack parameter**, and it specifies how ‘pessimistic’ the estimation of the average packet interarrival is, given a certain interarrival standard deviation. The computations involved in the above algorithm are quite simple to consist a significant overhead. If, however, more efficient calculations are required, the variance estimator can be replaced with the simpler to calculate mean prediction error, as it was done in [12] for the round-trip delay estimation in TCP implementations.

It is clear that the performance of HIP can depend strongly on the parameters α , β , γ , ϕ , and P_U^{MAX} . If they are chosen too conservatively, polling will seldom be used and the results will be similar to those of exclusive interrupts. At the other extreme of parameter selection, polling will be used too often and the overhead of unsuccessful polls will dominate, or the polling period will be too large, and the receive latencies will be increased without significant overhead reduction. In Section 6 we present some experimental results on the effect of these parameters.

A concern about the proposed algorithm is that the fractal (or self-similar) nature of data traffic [13] may deteriorate the predictability of the packet interarrivals. It is exactly because of this concern that we used real local-area network traffic traces in evaluating HIP. We observed through experimentation that frequently the simple coefficient-of-variation metric fails to accurately track the predictability of the arriving stream. This causes unsuccessful polls, or increased receive latencies. The constraint on the maximum number of unsuccessful polls serves as a protection mechanism, forcing HIP back to interrupt mode when the polling period is badly set over a sequence of packets.

5 Related Work

Several others have examined the trade-offs between interrupts and polling and have suggested hybrid

schemes. An interesting mechanism is the *Clocked Interrupts* [14, 15]. In that scheme, a fine-granularity timer determines the rate of interrupt events. A typical frequency range for this timer is 500Hz to 4KHz. Upon the expiration of the timer the network interface is polled for the arrival of one or more packets. If a packet has arrived, the interrupt service routine for the network interface is called to move the packets into the main memory. The receive overhead is reduced because the interrupt is not caused by the asynchronous event of a packet arrival, but by the timer expiration. This makes clocked interrupts quite similar with polling. Note however that this mechanism requires an additional fine-granularity high-frequency hardware timer. A second difference is that the clocked interrupts timer has a constant period that does not depend on the arriving network traffic, while HIP determines the polling period based on the measured packet interarrivals. This adaptation allows HIP to control the overhead-latency trade-off between interrupts and polling depending on the network interface load.

Another relevant scheme was presented in [5]. The authors focused on an effect called *Receive Livelock*, that occurs under heavy network load conditions in workstations that are configured as routers, firewalls, file servers, or promiscuous network monitors. Specifically, if packets arrive very frequently, the system can practically spend all the CPU time at the interrupt handling routine, while lower priority processes (including user processes) do not get the chance to process the packets to completion [citeqie:01]. This leads to extensive packet drops at intermediate buffering points, and thus, to a livelock situation. The authors of [5] implemented a mechanism where interrupts are only used at low network load conditions, while in high loads the interrupts are disabled and a polling thread is scheduled for reading the network interface. Every time a poll is executed, a certain *packet quota* is specified, i.e., the maximum number of packets that can be read in that poll. The quota is used for fairness purposes, when other tasks must also be permitted to make progress, so that to avoid the above livelock condition. If at the end of the polling some packets still remain at the NIC, the polling thread is executed again after a few milliseconds. Otherwise, the system switches back to interrupts. In principle, this mechanism is very similar to HIP. However, in HIP, the polling period is adjusted based on the observed packet interarrivals.

Aron and Druschel designed a *soft timers* OS facility that allows efficient scheduling of software events at microsecond granularity [10]. The basic idea behind soft timers is to take advantage of certain states in the execution of a system where an event handler can be invoked at low cost. Such states include the entry points of various kernel handlers, such as system calls and exception handlers. A drawback of soft timers is that they can only schedule events probabilistically. [16] has demonstrated, however, that under practical workloads it is possible to schedule events at intervals down to a few tens of microseconds, with rare delays up to a few hundred of microseconds. The soft timers facility can be combined with HIP in the following way: instead of scheduling the HIP polling thread at the granularity of the clock interrupt period, a soft timer can be scheduled to perform it instead. The execution frequency of the polling thread will then be adjusted by HIP, as described in the previous section.

In the context of message-passing parallel systems, a scheme that is similar to HIP is the *Polling Watchdog* [11]. In parallel systems the network latencies are part of the computational delays, and so it is critical to minimize them. The Polling Watchdog is a hardware extension at the NIC that limits the generation of interrupts to the cases where explicit polling fails to handle the packets quickly. The basic idea is that when a packet arrives at the NIC, a timer starts counting. If the packet is not removed from the NIC through polling within a given amount of time (the watchdog timeout period T_{wdog}), the watchdog interrupts the CPU. T_{wdog} is set to around 50 μ s, in order to strictly limit the maximum latency. In the EARTH-MANNA multiprocessor system¹ on which this scheme has been implemented, the cost of an interrupt is 4.5 μ s, and the cost of a poll is 400 ns. The major difference between the Polling Watchdog and HIP is that in the former the CPU always polls the NIC in every context switch (on the average every 50 μ s) and interrupts are used only to bound the maximum receive latency, while HIP switches between interrupts and variable-period polling, depending on the observed arriving network stream.

Another scheme that combines interrupts with polling in the context of message-passing parallel systems is used in the CM-5 [8]. For that system the time per poll is 1.6 μ s (0.6 to poll the interface and

1.0 to check the type of message and move it to memory), the interrupt overhead is 19 μ s, and the cost of enabling or disabling interrupts is about 4.3 μ s. The basic idea of the CM-5 scheme is to keep polling for incoming packets while servicing an interrupt for a packet arrival. The interrupt service routine exits only if there are no more packets waiting at the NIC. This is also called *batched interrupts* and it is a common practice in current network interface drivers [5]. The virtue of this scheme is that if the average packet interarrival is much smaller than the interrupt overhead, then multiple packets can be serviced in a single interrupt. In multiprocessor systems, the network bandwidth is very high and, therefore, the average packet interarrival times are often in the range of a few microseconds. Hence, batched interrupts are quite effective. For more conventional workstations and networks, not many packets can arrive within a single interrupt service period. For example, for a 10Mbps Ethernet the minimum time between packet arrivals is about 51.2 μ s, and so in a 5 μ s interrupt handling period at most two packets can be captured. HIP does both batched interrupts and batched polls. In addition, it attempts to reduce the receive overhead when packets arrive every several hundreds or thousands of microseconds.

6 Evaluation

We wrote an event-driven simulator to quantitatively evaluate the overhead and latency that results from HIP, as well as from schemes that exclusively use interrupts or polling. The main goal of this study is to examine the relative overhead and latency that results from these different interface schemes, rather than to predict system-specific absolute performance measures. Additionally, we are interested to examine the behavior of the packet interarrival predictors using real network traces, and to find appropriate values for the HIP parameters.

The events that the simulator captures (and their symbols) are:

- Packet arrival (A)
- End of interrupt period (I)
- Start of polling period (P_s)
- End of polling period (P_f)

¹based on Intel i860 XP processors.




Figure 2: Two timelines with packet arrivals (A), interrupt period ends (I), polling period starts (P_s), and polling period ends (P_f).

The packet arrivals are determined from real LAN traffic traces (described later). In the interrupt mode, an interrupt period is initiated with a packet arrival, unless the packet arrival occurs inside an interrupt service period. In the latter case, we receive all the packets within the same interrupt (batched interrupts). The length of the interrupt period is equal to the interrupt overhead $V_I + V(B)$. The start of a polling event is initiated based on the polling period T_P , and its duration is determined based on the polling overhead $V_P + V(B)$. Batched polls are also supported. After every packet arrival, the predictions \bar{I} and σ_f^2 are updated and the next polling period is determined (if in Polling-mode). Statistics regarding the average receive overhead per packet and the average latency per packet are reported at the end of the simulation run. The timelines in Figure 2 illustrate the above events, and the interrupt and polling overheads.

The characteristics of the network traffic traces that we use in this paper appear in Table 1. All traces were collected using the ‘snoop’ packet filter, observing the packets that are destined to a Sun Ultra-1 workstation from a 10Mbps Ethernet segment. It is important to note that on this specific Ethernet segment there are several dozens of connected workstations, and our workstation frequently received ARP, NFS, DNS, and other kinds of packets that the user does not normally see. The traces are classified in three classes, depending on the average bit-rate that arrives at our host (high, medium, and low). The first class (T-high) consists of two traces that were generated while long FTP sessions with nearby workstations were ongoing. The achieved average throughput in these sessions is close to 3 Mbps, so these streams may represent the kind of high-bandwidth multimedia streams that computers will be receiving in the near future. The second class (T-medium) consists of two traces that were generated while Web-browsing

sessions were performed and correspond to an average incoming throughput of several tens of kbps. The third class (T-low) of traces were generated while the workstation was otherwise idle. In these traces the average incoming throughput is only a few kbps, mainly due to ARP, NFS, and DNS traffic.

The parameter values that we used in the simulations are shown in Table 2. The fixed parameters (V_I , V_P , $V(B)$, V_S) were selected based on reported values in the literature [16, 5]. Note these particular numbers can vary widely from system to system, depending on the I/O system design and implementation. The fact though that the polling overhead is almost an order of magnitude smaller than the interrupt overhead seems to be rather generally true. The per-byte overhead $V(B)$ was assumed to be quite low (500ns for 1500 byte packets), as would be the case with a DMA-capable NIC. The minimum polling period was set to 1ms, assuming a soft clock period of the same rate, while the maximum tolerable packet latency was set to 10ms. The HIP-related parameters (α , β , γ , ϕ , P_U^{MAX}) were tuned through extensive simulations with a wide range of traffic traces (not only the reported ones here).

The comparison of HIP with the interrupts-only and the polling-only network interface schemes is shown in Table 3. Note that, the simulated polling scheme uses the same polling-period adaptation algorithm as HIP, without ever switching to interrupt mode though. Conventional polling schemes, on the other hand, either busy-wait on the NIC, or use a constant polling period. The average overhead per packet and the average latency per packet are shown for each of these three network interface schemes. Several other performance measures of HIP, such as the average number of packets that are captured through interrupts and through polling, are also shown.

As expected, the interrupt overhead remains almost constant independent of the incoming traffic’s rate. On the contrary, the polling overhead keeps increasing as the incoming traffic’s rate decreases, due to many unsuccessful polls. HIP achieves the low overhead of polling at high input rates, while for lower input rates it consistently introduces a lower overhead than the interrupt-based and polling-based schemes. For the high rate traces, the average HIP overhead per packet is almost six times smaller than the interrupts-only overhead. On the negative side,

	T-High-1	T-High-2	T-Med-1	T-Med-2	T-Low-1	T-Low-2
Duration (sec)	11.5	16.7	255.0	341.9	1118.7	1302.3
Packets Received	4502	5407	3736	3997	3779	4044
Average Interarrivals (ms)	2.5	3.1	68.2	85.6	296.0	322.0
Avg. Packet Length (bytes)	904	949	421	561	90	101


Table 1: Characteristics of the network traces that were used in the simulations.

V_I	V_P	$V(B)$	V_S	T_P^{MIN}	T_P^{MAX}	α	β	γ	ϕ	P_U^{MAX}
$5\mu s$	$0.5\mu s$	$0.5\mu s$ B=1500	$1\mu s$	1ms	10ms	0.3	0.3	2	2.5	10


Table 2: Parameter values used in simulations.

	T-High-1	T-High-2	T-Med-1	T-Med-2	T-Low-1	T-Low-2
Interrupts: Avg-Overhead (μs)	5.3	5.3	5.1	5.2	5.0	5.0
Polling: Avg-Overhead (μs)	0.9	0.9	6.3	5.4	18.4	20.3
HIP: Avg-Overhead (μs)	0.9	0.9	2.6	4.1	3.3	3.4
Interrupts: Avg-Delay (μs)	5.3	5.3	5.1	5.2	5.0	5.0
Polling: Avg-Delay (μs)	1698	1640	3455	4598	3919	3975
HIP: Avg-Delay (μs)	1557	1547	1074	1321	393	367
HIP: Total # of Interrupts	103	111	1750	2792	2329	2531
HIP: Avg. Packets per Interrupt	1.0	1.0	1.0	1.0	1.0	1.0
HIP: Total # of Polls	4415	5458	2429	2726	1404	1504
HIP: Total # of Unsucc. Polls	1813	2224	1511	2093	825	869
HIP: Avg. Packets per Poll	1.0	1.0	0.8	0.4	1.0	1.0


Table 3: Comparison of HIP with the schemes that are based exclusively on interrupts and polling.




(a) The effect of γ on the average overhead per packet



(b) The effect of γ on the average delay per packet



(a) The effect of ϕ on the average overhead per packet



(b) The effect of ϕ on the average delay per packet

Figure 3: The average per packet overhead and delay, as a function of the parameter γ .

Figure 4: The average per packet overhead and delay, as a function of the parameter ϕ .

HIP introduces a considerably larger average per packet latency (between $400\mu s$ and $1.5ms$) compared to the interrupt-based scheme. Note that the average delay using polling-only can be much larger than that of HIP, especially for low-rate streams. The last five lines of Table 3 show the mix between interrupts and polls for HIP. The high incoming rate traces are captured mainly through polling, while the low rate traces lead to most interrupts and many unsuccessful polls. In general, the large number of unsuccessful polls relatively to the total number of polls can be attributed to interarrival prediction failures because of the traffic burstiness.

Figure 3 shows the effect of the parameter γ on the average packet overhead and latency. It is clear that as γ increases, we expect less regularity in the incoming stream and polling is selected more frequently. This results in decreasing overhead and increasing latency. A value of γ between 2 and 3 seems to achieve a good trade-off between average overhead and latency.

Figure 4 shows the effect of the parameter ϕ on the average packet overhead and latency. As ϕ increases, the polling period also increases. For high-rate incoming streams this reduces the average overhead, since more packets are received in each poll and less unsuccessful polls are performed. This effect is less important in low rate streams. On the other hand, larger values of ϕ can increase the average latency significantly. We found that a value of ϕ between 2 and 3 is a good middle point.

7 Conclusions

We proposed a hybrid interrupt-polling (HIP) scheme for the network interface. HIP exploits the trade-off between decreased receive-overhead and increased receive-latency. The careful selection of the related parameters allows the system designer to set this trade-off to the appropriate operating point. Through trace-driven simulations, we showed that HIP is more effective when the network workload consists of high-bandwidth streams, such as multimedia traffic. In this paper we focused on the underlying ideas and the general architecture, rather than on implementation specific problems. We expect, though, that several challenges can arise at that phase. Identifying and addressing those issues is something that we plan to do in the future.

References



- [1] D. Banks and M. Prudence, "A high-performance network architecture for a PA-RISC workstation," *IEEE Journal on Selected Areas in Communications*, vol. 11, pp. 191–202, Feb. 1993.
- [2] C. B. S. Traw and J. M. Smith, "Hardware-software organization of a high-performance ATM host interface," *IEEE Journal on Selected Areas in Communications*, vol. 11, pp. 240–253, Feb. 1993.
- [3] B. Davie, "The architecture and implementation of a high-speed host interface," *IEEE Journal*

- on *Selected Areas in Communications*, vol. 11, pp. 228–239, Feb. 1993.
- [4] K. K. Ramakrishnan, "Performance considerations in designing network interfaces," *IEEE Journal on Selected Areas in Communications*, vol. 11, pp. 203–219, Feb. 1993.
 - [5] J. C. Mogul and K. K. Ramakrishnan, "Eliminating receive livelock in an interrupt-driven kernel," *ACM Transactions on Computer Systems*, vol. 15, pp. 217–252, Aug. 1997.
 - [6] S. S. Mukherjee and M. D. Hill, "A survey of user-level network interfaces for system area networks," Tech. Rep. TR 1340, Computer Sciences Department, University of Wisconsin- Madison, Feb. 1997.
 - [7] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, *RTP: A Transport Protocol for Real-Time Applications*, Jan. 1996. RFC 1889.
 - [8] D. A. Patterson and J. L. Hennessy, *Computer Architecture, A Quantitative Approach*. Morgan Kaufmann, 1996.
 - [9] U. Vahalia, *UNIX Internals, the new frontiers*. Prentice Hall, 1996.
 - [10] M. Aron and P. Druschel, "Soft Timers: Efficient Microsecond Software Timer Support for Network Processing," *ACM Transactions on Computer Systems*, vol. 18, pp. 197–228, Aug. 2000.
 - [11] O. Maquelin, G. R. Gao, H. H. J. Hum, K. B. Theobald, and X. Tian, "Polling watchdog: Combining polling and interrupts for efficient message handling," in *Proceedings of the 23rd Annual International Symposium on Computer Architecture*, pp. 179–188, 1996.
 - [12] V. Jacobson, "Congestion Avoidance and Control," in *Proceedings of ACM SIGCOMM*, pp. 314–329, Sept. 1988.
 - [13] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," in *Proceedings of ACM SIGCOMM*, pp. 100–113, Sept. 1995.
 - [14] J. D. Chung, C. B. S. Traw, and J. M. Smith, "Event-signaling within higher performance network subsystems," in *Proceedings, High Performance Communications Subsystems*, pp. 220–225, Aug. 1995.
 - [15] J. M. Smith and C. B. S. Traw, "Operating systems support for end-to-end Gbps networking," *IEEE Network*, vol. 7, pp. 44–52, Feb. 1993.
 - [16] M. Aron and P. Druschel, "Soft Timers: Efficient Microsecond Software Timer Support for Network Processing," in *Proceedings of ACM SOSP*, Dec. 1999.

U. S. DEPARTMENT OF COMMERCE
COMMISSIONER FOR PATENTS
P.O. BOX 1450
ALEXANDRIA, VA 22313-1450
IF UNDELIVERABLE RETURN IN TEN DAYS

OFFICIAL BUSINESS

AN EQUAL OPPORTUNITY EMPLOYER



BEST AVAILABLE COPY

